A Universal Decoder Relative to a Given Family of Metrics

Nir Elkayam and Meir Feder
Department of Electrical Engineering - Systems
Tel-Aviv University, Israel
Email: nirelkayam@post.tau.ac.il, meir@eng.tau.ac.il

arXiv:1401.6432v3 [cs.IT] 27 Apr 2014

Abstract

Consider the following framework of universal decoding suggested in [1]. Given a family of decoding metrics and random coding distribution (prior), a single, universal, decoder is optimal if for any possible channel the average error probability when using this decoder is better than the error probability attained by the best decoder in the family up to a subexponential multiplicative factor. We describe a general universal decoder in this framework. The penalty for using this universal decoder is computed. The universal metric is constructed as follows. For each metric, a canonical *metric is defined and conditions for the given prior to be* normal *are given. A sub-exponential set of canonical metrics of normal prior can be merged to a single universal optimal metric. We provide an example where this decoder is optimal while the decoder of [1] is not.*

I. Introduction

Metric based decoders are usually used to decode codes in digital communication. Typically, the metric tries to capture the most likely codeword that was transmitted. When the channel over which the codeword was transmitted is known at the receiver, the *Maximum Likelihood* (ML) decoder is used for optimal performance (in terms of block error rate). When the channel is not known at the receiver, other solutions are required.

Usually another assumption is made, for example, the channel $W_\theta(\mathbf{y}|\mathbf{x})$ belongs to some family of channels indexed by $\theta \in \Theta$. The family of channels can have some measure defined over it (Bayesian approach), *e.g.*, fading channel, or the channel is one of several possible channels (deterministic approach), *e.g.*, discrete memoryless channels (DMCs). The performance in the first case can be measured over the channel realization, and in the latter per specific channel.

A universal decoding is *optimal* with respect to a given family of channels if the performance when using the universal decoder is not worse (in the error exponent sense) than when the optimal decoder is used; specifically, the error exponent is not smaller than that of the optimal decoder.

In [3], Goppa offered the *maximum mutual information* (MMI) decoder, which decides in favor of the codeword having the maximum empirical mutual information with the channel output sequence. Other universal decoders have been suggested over the years for ISI channels [4], [5], [6],Interference channels [7], finite state channels [8], [9], Individual channels [10], and [11]. In [12], Feder and Lapidoth provide a fairly general construction of a universal decoder for a given class of channels. Their construction has two key points. The first is the use of uniform random coding distribution (prior), which means that all codewords have the same probability; this allows the use of the *merged decoder*. The second is the definition of a separable family, which means that there is a countable set of channels that can be used as representative, in the sense that if the performance of the universal decoder is good over these channels, then it would be good over the whole family of channels.

In [1] Merhav proposed another framework of universal decoding, namely, universality relative to a given set of decoding metric over arbitrary channels, where optimality means that the performance of the universal decoder is not worse (in the error exponent sense) than the performance when any other metric from the family is used. The decoder proposed by Merhav in [1] uses an equivalent of "types" that is a set of input and output sequences of the channel for which the metrics of the family provides the same value. The universal decoder that is based on ties in the metric, thus, does not seem to be the most general.

In this paper we generalize the construction in [1], which in some sense also generalizes Feder and Lapidoth's universal decoder, [12]. The quantity $p_{e,m,\mathbf{x},\mathbf{y}}$, which is the error probability when using the metric m, given that $\mathbf{x}$ was transmitted and $\mathbf{y}$ was received, is defined. This quantity plays an important role in the computation of the average error probability and can be used to compare the performance between different metrics. Based on $p_{e,m,\mathbf{x},\mathbf{y}}$ we will define, for a metric m, a "canonical metric" $\overline{m}(\mathbf{x},\mathbf{y}) = -\log p_{e,m,\mathbf{x},\mathbf{y}}$, which is equivalent to m. The canonical metrics (or the associated error probabilities $p_{e,m,\mathbf{x},\mathbf{y}}$) serve, in turn, to define a universal decoder relative to a set of metrics, called the **Generalized Minimum Error Test**. We derive the penalty of using this universal metric, and show when this metric is optimal (in the error exponent sense).

The definition of the canonical metric is related to the definition in [13] of *empirical rate function*. This rate function, which is a function of the transmitted word $\mathbf{x}$ and the received word $\mathbf{y}$, was defined in an attempt to remove the probabilistic

This work was partially supported by the Israeli Science Foundation (ISF), grant no. 634/09.

assumption in communication and get a rate function that fits the individual sequences of input and output. Specifically, after we introduce the notion of normal prior, we show that for every metric m over a normal prior there exists a "tight" empirical rate function, *i.e.*, the canonical metric, and $\overline{m}(\mathbf{x},\mathbf{y}) \approx \log\left(\frac{P(\mathbf{x}|\mathbf{y})}{Q(\mathbf{x})}\right)$. Here $\approx$ means equal up to a sub-exponential factor and $P(\mathbf{x}|\mathbf{y})$ is a probability distribution over $\mathcal{X}$. A set of countable "tight" rate functions can be merged into a single universal decoding metric that is optimal with respect to all the decoding metrics in the set. This is the generalization of the merged decoder of [12].

II. Notation

This paper uses bold lower case letters (*e.g.*, **x**) to denote a particular value of the corresponding random variable denoted in capital letters (*e.g.*, **X**). Calligraphic fonts (*e.g.*, $\mathcal{X}$) represent a set. Throughout this paper $\log$ will be defined as base 2 unless otherwise indicated. $Pr\{A\}$ will denote the probability of the event A. We will use the o notation. Specifically, $o(n)$ will be used to indicate a sequence of the form: $n \cdot \lambda_n$ where $\lambda_n \underset{n\to\infty}{\to} 0$ and $o(1)$ is a sequence that goes to 0. A sequence of positive numbers χ_n is said to be **subexponential** if $\lim_{n\to\infty} \frac{1}{n}\log(\chi_n) = 0$. With the o notation this is just $\chi_n = 2^{o(n)}$

III. Preliminaries

A rate-R blocklength-n, **random code** with prior $Q_n(\mathbf{x})$ is:

$$\mathcal{C} = \left\{\mathbf{x}(1), \mathbf{x}(2), ..., \mathbf{x}(2^{nR})\right\} \subset \mathcal{X}^n, \tag{1}$$

where each codeword $\mathbf{x}(i)$ is drawn independently according to the distribution $Q_n(\mathbf{x})$. The **encoder** maps the input message k, drawn uniformly over the set $\left\{1, ..., 2^{nR}\right\}$, to the appropriate codeword, *i.e.*, $\mathbf{x}(k)$. A **decoding metric** is a function $m : \mathcal{X}^n \times \mathcal{Y}^n \to \mathbb{R}$. The **decoder** associated with the decoding metric m is the function $\mathcal{D}_m : \mathcal{Y}^n \to \left\{1, ..., 2^{nR}\right\}$ defined by:

$$\mathcal{D}_m(\mathbf{y}) = \hat{i} = argmax_{1\le i\le 2^{nR}} \left(m(\mathbf{x}(i), \mathbf{y})\right). \tag{2}$$

Definition 1. The **average error probability** associated with the decoder $\mathcal{D}_m$ over the channel $W_n(\mathbf{y}|\mathbf{x})$ is denoted by $\bar{P}_{e,m,W}(R)$. The error probability captures the randomness in the codebook selection, the transmitted message, and the channel output.

Remark 1. We will generally refer to a metric as a decoder where it would be understood that we refer to the decoder associated with the given metric.

Remark 2. When the decoding metric $m(\mathbf{x},\mathbf{y}) = W_n(\mathbf{y}|\mathbf{x})$, the error probability is minimized and the decoder is called the Maximum Likelihood decoder. The **mismatch** case is when the metric is not matched to the channel and may result in performance degradation.

Definition 2. For a family of decoding metrics:

$$\mathcal{M}_n = \{m_\theta(\mathbf{x},\mathbf{y}) : \theta \in \Theta_n, \mathbf{x} \in \mathcal{X}^n, \mathbf{y} \in \mathcal{Y}^n\}. \tag{3}$$

Denote:

$$\bar{P}_{e,\mathcal{M}_n,W_n}(R) = \min_{\theta\in\Theta_n} \bar{P}_{e,m_\theta,W_n}(R), \tag{4}$$

the minimum error probability over the channel $W_n(\mathbf{y}|\mathbf{x})$ when the best metric matched for this channel is used.

Definition 3. A sequence of decoding metrics u_n, independent of θ, and their associated decoders is called **universal**. The sequence of universal decoding metrics u_n (and their decoders) is **optimal** (in the error exponent sense) with respect to the family $\mathcal{M}_n$ if:

$$\lim_{n\to\infty} \frac{1}{n} \log\left(\frac{\bar{P}_{e,u_n,W(\mathbf{y}|\mathbf{x})}(R)}{\bar{P}_{e,\mathcal{M}_n,W(\mathbf{y}|\mathbf{x})}(R)}\right) \le 0$$

for every channel $W(\mathbf{y}|\mathbf{x})$.

Consider a metric $m(\mathbf{x},\mathbf{y})$. The pairwise error given that **x** was transmitted and **y** was received is given by:

$$p_{e,m,\mathbf{x},\mathbf{y}} = Q_n(m(\mathbf{X},\mathbf{y}) \ge m(\mathbf{x},\mathbf{y})). \tag{5}$$

The probability given by (5) is for a single (random) codeword **X**, to beat the transmitted word **x** given that the received word was **y**. Notice that in case of a tie, we are assuming that an error has occurred.

The term $p_{e,m,\mathbf{x},\mathbf{y}}$ can be used to compute the average error probability for any rate R as can be seen in the following lemma.

Lemma 1.

$$\frac{1}{2} \leq \frac{\bar{P}_{e,m,W}(R)}{\mathbb{E}\left(\min\left(1, 2^{nR} \cdot p_{e,m,\boldsymbol{x},\boldsymbol{y}}\right)\right)} \leq 1, \tag{6}$$

where the expectation is with respect to $Q(\boldsymbol{x})W(\boldsymbol{y}|\boldsymbol{x})$.

This lemma is given in [1]. The upper bound follows from the union bound. The lower bound follows from a lemma proved by Shulman in [14, Lemma A.2]; namely, that the union clipped to unity is tight up to factor 2 (for pairwise independent events).

Definition 4. We say that the metric u_n **dominates** m_n if for all $\mathbf{x}, \mathbf{y}$:

$$p_{e,u_n,\mathbf{x},\mathbf{y}} \leq p_{e,m_n,\mathbf{x},\mathbf{y}} \cdot 2^{n \cdot \lambda_n} \tag{7}$$

with with $\lambda_n \underset{n\to\infty}{\to} 0$.

The following lemma, which also appears in [1], follows from lemma (1):

Lemma 2.

$$\bar{P}_{e,u_n,W}(R) \leq 2 \cdot 2^{n \cdot \lambda_n} \cdot \bar{P}_{e,m_n,W}(R) \tag{8}$$

$$\bar{P}_{e,u_n,W}(R) \leq 2 \cdot \bar{P}_{e,m_n,W}(R + \lambda_n). \tag{9}$$

In particular, if metric u_n dominates m_n then the average error exponent when using the metric u_n is not worse than using the metric m_n.

Obviously, an optimal universal decoding metric is a metric that dominates all the metrics in the family (3).

IV. The Generalized minimum error test

In this section we propose our universal decoder, the **Generalized Minimum Error Test (GMET)**. The decoder estimates the metric that minimizes the pairwise error probability.

Definition 5. For the given family of metrics (3), let:

$$2^{-n \cdot U_n(\mathbf{x},\mathbf{y})} \triangleq \min_{\theta \in \Theta_n} p_{e,m_\theta,\mathbf{x},\mathbf{y}}. \tag{10}$$

$U_n(\mathbf{x}, \mathbf{y})$ is the GMET. Let:

$$K_n(\Theta_n) \triangleq \max_{\mathbf{y}} \mathbb{E}_{Q_n(\mathbf{x})}\left(2^{n \cdot U_n(\mathbf{X},\mathbf{y})}\right). \tag{11}$$

$K_n(\Theta_n)$ is termed the **redundancy** associated with the decoder.

The performance degradation of the proposed universal decoder relative to any metric $m_\theta(\mathbf{x}, \mathbf{y})$ in the family can be measured through the redundancy as demonstrated in the following Theorem.

Theorem 1. *For any $\theta \in \Theta_n$*

$$p_{e,U_n,\boldsymbol{x},\boldsymbol{y}} \leq p_{e,m_\theta,\boldsymbol{x},\boldsymbol{y}} \cdot K_n(\Theta_n). \tag{12}$$

In particular, if $K_n(\Theta_n)$ is sub-exponential then the universal decoder $U_n(\boldsymbol{x},\boldsymbol{y})$ is optimal.

Proof:

$$\begin{aligned} p_{e,U_n,\mathbf{x},\mathbf{y}} &= Q_n\left(U_n(\mathbf{X},\mathbf{y}) \geq U_n(\mathbf{x},\mathbf{y})\right) \\ &= Q_n\left(2^{n \cdot U_n(\mathbf{X},\mathbf{y})} \geq 2^{n \cdot U_n(\mathbf{x},\mathbf{y})}\right) \\ &\overset{(a)}{\leq} \mathbb{E}_{Q_n(\mathbf{x})}\left(2^{n \cdot U_n(\mathbf{X},\mathbf{y})}\right) \cdot 2^{-n \cdot U_n(\mathbf{x},\mathbf{y})} \\ &\overset{(b)}{\leq} p_{e,m_\theta,\mathbf{x},\mathbf{y}} \cdot K_n(\Theta_n) \end{aligned}$$

where (a) follows from Markov's inequality and (b) from the definition of $K_n(\Theta_n)$ (11) and $U_n(\mathbf{x}, \mathbf{y})$ (10). If $K_n(\Theta_n)$ is sub-exponential then $U_n(\mathbf{x}, \mathbf{y})$ dominates every metric m_θ. By Lemma 2 the universal decoder is optimal. ■

Theorem 1 provides a quite general construction of a universal decoder. However, bounding $K_n(\Theta_n)$ might not be an easy task.

A. Conditions for universal decoding

Obviously, for a family of a *single* metric there must be a universal decoder, *i.e.*, the same single metric or an equivalent metric. By analyzing this simple case we get a core understanding of when the proposed GMET decoder is optimal.

Definition 6. For a given matric $m(\mathbf{x},\mathbf{y})$ let:

$$\bar{m}(\mathbf{x},\mathbf{y}) \triangleq -\frac{1}{n}\log\left(p_{e,m,\mathbf{x},\mathbf{y}}\right). \tag{13}$$

$\bar{m}(\mathbf{x},\mathbf{y})$ is called **the canonical metric** associated with $m(\mathbf{x},\mathbf{y})$.

The canonical metric $\bar{m}$, which is equivalent to the original metric m as it induces the same order of the candidate words, is the proposed universal decoder for the family of a single metric $m(\mathbf{x},\mathbf{y})$. The next definition captures the cases where this universal metric is optimal in the sense of Theorem 1.

Definition 7.

1) The metric $m(\mathbf{x},\mathbf{y})$ over the prior $Q_n(\mathbf{x})$ is **normal** if
$$\mathbb{E}_{Q_n(\mathbf{x})}\left(2^{n\cdot\bar{m}(\mathbf{X},\mathbf{y})}\right) \leq 2^{n\cdot\lambda_n} \tag{14}$$
uniformly over $\mathbf{y}$, where $\lambda_n \underset{n\to\infty}{\to} 0$.
2) The prior $Q_n(\mathbf{x})$ is normal if each metric over $Q_n(\mathbf{x})$ is normal (uniformly over all the metrics, *i.e.*, the bound in (14) is independent of the metric).

Normal metrics are exactly those metrics that might admit an optimal universal decoder.

The following lemma provides a sufficient condition for a given discrete prior $Q_n(\mathbf{x})$ to be normal. The condition is very mild and easy to check.

Lemma 3. *Let $\chi_n = -\log\left(\min_{\boldsymbol{x}:Q_n(\boldsymbol{x})\neq 0} Q_n(\boldsymbol{x})\right)$. If χ_n is sub-exponential, then $Q_n(\mathbf{x})$ is normal.*

Proof: The proof is given in the appendix (Proposition 5) where we also introduce the concept of "rate function" and relate it to the canonical metric. ■

B. The Merged decoder

If the family of metrics (3) is finite, we can "merge" all the metrics and get a bound on $K_n(\Theta_n)$. If the family size grows sub-exponentially over the normal prior, then the family admits a universal decoder.

Theorem 2. *Suppose $Q_n(\boldsymbol{x})$ is normal and Θ_n is finite, then:*
1) $K_n(\Theta_n) \leq |\Theta_n| \cdot 2^{o(n)}$
2) *If $|\Theta_n|$ grows sub-exponentially, then the family Θ_n admits an optimal universal decoder.*

Proof: For each $\mathbf{y}$:

$$\begin{aligned}\mathbb{E}_{Q_n(\mathbf{x})}\left(2^{n\cdot U_n(\mathbf{X},\mathbf{y})}\right) &= \mathbb{E}_{Q_n(\mathbf{x})}\left(\frac{1}{\min_{\theta\in\Theta_n} p_{e,m_\theta,\mathbf{x},\mathbf{y}}}\right)\\ &\overset{(a)}{\leq} \sum_{\theta\in\Theta_n}\mathbb{E}_{Q_n(\mathbf{x})}\left(\frac{1}{p_{e,m_\theta,\mathbf{x},\mathbf{y}}}\right)\\ &\overset{(b)}{\leq} |\Theta_n|\cdot 2^{o(n)}\end{aligned}$$

(a) follows by taking $\max$ over all $\mathbf{y}$ and (b) from Theorem (1). ■

C. Approximation of the universal decoder

As noted also in [1], the proof actually gives conditions for a universal decoding metric to be optimal. Suppose that there exists another universal metric $U'_n(\mathbf{x},\mathbf{y})$ such that:

$$2^{-n\cdot U'_n(\mathbf{x},\mathbf{y})} \leq 2^{-n\cdot U_n(\mathbf{x},\mathbf{y})+o(n)} \tag{15}$$

and

$$\mathbb{E}_{Q_n(\mathbf{x})}\left(2^{n\cdot U'_n(\mathbf{x},\mathbf{y})}\right) \leq 2^{o(n)}. \tag{16}$$

Then it follows that:

$$\begin{aligned} p_{e,U'_n,\mathbf{x},\mathbf{y}} &= Q_n\left(U'_n(\mathbf{X},\mathbf{y}) \geq U'_n(\mathbf{x},\mathbf{y})\right) \\ &\overset{(a)}{\leq} \mathbb{E}_{Q_n(\mathbf{x})}\left(2^{n\cdot U'_n(\mathbf{X},\mathbf{y})}\right)\cdot 2^{-n\cdot U'_n(\mathbf{x},\mathbf{y})} \\ &\overset{(b)}{\leq} 2^{o(n)}\cdot 2^{-n\cdot U_n(\mathbf{x},\mathbf{y})} \\ &\overset{(c)}{\leq} p_{e,m_\theta,\mathbf{x},\mathbf{y}}\cdot K_n(\Theta_n)\cdot 2^{o(n)} \end{aligned}$$

where (a) follows from Markov's inequality, (b) is the given conditions, and (c) is Theorem 1. In particular, if $K_n(\Theta_n)$ is subexponential, then $U'_n(\mathbf{X},\mathbf{y})$ is also optimal. Now,

$$\begin{aligned} p_{e,m_\theta,\mathbf{x},\mathbf{y}} &= Q\left(m_\theta(\mathbf{X},\mathbf{y}) \geq m_\theta(\mathbf{x},\mathbf{y})\right) \\ &\geq Q\left(m_\theta(\mathbf{X},\mathbf{y}) = m_\theta(\mathbf{x},\mathbf{y})\right) \\ &\geq Q\left(\forall\tau\in\Theta_n, m_\tau(\mathbf{X},\mathbf{y}) = m_\tau(\mathbf{x},\mathbf{y})\right) \end{aligned}$$

So there are several general approximations that might be used instead of the original proposed decoder, *i.e.*:

$$2^{-n\cdot U_{n,1}(\mathbf{x},\mathbf{y})} = \min_{\theta\in\Theta_n} Q\left(m_\theta(\mathbf{X},\mathbf{y}) = m_\theta(\mathbf{x},\mathbf{y})\right). \tag{17}$$

and

$$2^{-n\cdot U_{n,2}(\mathbf{x},\mathbf{y})} = Q\left(\forall\tau\in\Theta_n, m_\tau(\mathbf{X},\mathbf{y}) = m_\tau(\mathbf{x},\mathbf{y})\right), \tag{18}$$

where in the definition of $U_{n,2}(\mathbf{x},\mathbf{y})$ we can avoid the minimization over θ as it is effectively done in the probability computation. For each of these decoders the condition (16) should be checked and universality(optimal) of the approximation implies universality of the GMET decoder, but not vice versa (*c.f.* example V-C). Notice that $U_{n,2}(\mathbf{x},\mathbf{y})$ is the decoder proposed by Merhav in [1]. These approximations might be useful when computation of the exact GMET is hard and these approximation are easier to calculate.

Remark 3. The decoder $U_{n,2}(\mathbf{x},\mathbf{y})$ has a nice property of being a "asymptotically tight rate function" (when it is optimal) in the following sense: The use of the Markov's inequality gives the upper bound on $p_{e,U_{n,2},\mathbf{x},\mathbf{y}}$:

$$p_{e,U_{n,2},\mathbf{x},\mathbf{y}} \leq 2^{-n\cdot U_{n,2}(\mathbf{x},\mathbf{y})}\cdot 2^{o(n)}.$$

However, for $U_{n,2}$ it is also true that $p_{e,U_{n,2},\mathbf{x},\mathbf{y}} \geq 2^{-n\cdot U_{n,2}(\mathbf{x},\mathbf{y})}$. To see this, recall from [1, Eq. (4),(5)]:

$$\mathcal{T}_n(\mathbf{x}|\mathbf{y}) = \{\mathbf{x}'\in\mathcal{X}^n : \forall\theta\in\Theta_n, m_\theta(\mathbf{x}',\mathbf{y}) = m_\theta(\mathbf{x},\mathbf{y})\}$$

and

$$2^{-n\cdot U_{n,2}(\mathbf{x},\mathbf{y})} = Q_n\left(\mathcal{T}_n(\mathbf{x}|\mathbf{y})\right).$$

Then:

$$\begin{aligned} p_{e,U_{n,2},\mathbf{x},\mathbf{y}} &= \sum_{\mathbf{x}':U_{n,2}(\mathbf{x}',\mathbf{y})\geq U_{n,2}(\mathbf{x},\mathbf{y})} Q_n(\mathbf{x}') \\ &\geq \sum_{\mathbf{x}':U_{n,2}(\mathbf{x}',\mathbf{y})= U_{n,2}(\mathbf{x},\mathbf{y})} Q_n(\mathbf{x}') \\ &\overset{(a)}{\geq} \sum_{\mathbf{x}'\in\mathcal{T}_n(\mathbf{x}|\mathbf{y})} Q_n(\mathbf{x}') \\ &= Q_n\left(\mathcal{T}_n(\mathbf{x}|\mathbf{y})\right) \\ &= 2^{-n\cdot U_{n,2}(\mathbf{x},\mathbf{y})} \end{aligned}$$

where (a) follows because $U_{n,2}(\mathbf{x},\mathbf{y}) = U_{n,2}(\mathbf{x}',\mathbf{y})$ for $\mathbf{x}'\in\mathcal{T}_n(\mathbf{x}|\mathbf{y})$. Putting these together we have:

$$-\frac{1}{n}\log\left(p_{e,U_{n,2},\mathbf{x},\mathbf{y}}\right) = U_{n,2}(\mathbf{x},\mathbf{y}) + o(1) \tag{19}$$

So we see that the metric $U_{n,2}$ is almost in the canonical form (up to a vanishing term), which means that we can use it to evaluate the performance of the decoder (instead of using $p_{e,U_{n,2},\mathbf{x},\mathbf{y}}$). It is not apparent that U_n in general has the "asymptotically

tight" property of $U_{n,2}$. This means that $U_n(\mathbf{x},\mathbf{y})$ as a rate function provides only the lower bound on performance, and it might be that the performance is better in the sense that $-\frac{1}{n}\log\left(p_{e,U_n,\mathbf{x},\mathbf{y}}\right)$ might be larger than $U_n(\mathbf{x},\mathbf{y})$ by a non-vanishing term.

V. Examples

A. Normal priors examples

Example 1. *Let $Q_n(\boldsymbol{x}^n)$ be an i.i.d. probability distribution function, namely, $Q_n(\boldsymbol{x}^n)=\prod_{i=1}^n Q(\boldsymbol{x}_i)$. Obviously,*

$$\min_{\boldsymbol{x}:Q_n(\boldsymbol{x}^n)\neq 0} Q_n(\boldsymbol{x}^n)=\left(\min_{\boldsymbol{x}:Q(\boldsymbol{x})\neq 0} Q(\boldsymbol{x})\right)^n$$

and

$$-\log\left(\min_{\boldsymbol{x}:Q_n(\boldsymbol{x}^n)\neq 0} Q_n(\boldsymbol{x}^n)\right)=n\cdot\log\left(\min_{\boldsymbol{x}:Q(\boldsymbol{x})\neq 0} Q(\boldsymbol{x})\right),$$

which is subexponential (actually, linear in n).

Example 2. *Let $Q_n(\boldsymbol{x})$ be uniform over some set $B_n\subset\mathcal{X}^n$. If $|\mathcal{X}|$ is finite, then $|B_n|\leq|\mathcal{X}|^n$ and $-\log\left(|B_n|^{-1}\right)\leq-\log\left(|\mathcal{X}|^{-n}\right)=n\cdot\log\left(|\mathcal{X}|\right)$, which is subexponential. Hence, $Q_n(\boldsymbol{x})$ is normal.*

B. DMC and Constant type metrics

It is instructive to validate the fact that the family of metrics for DMCs, and more generally, metrics that are constant on types admits universal decoding over any normal prior. To see this notice that at first glance the metric family Θ_n is infinite, but since only the order on the input space induced by the metric matters, there are only a finite number of metrics that the minimum in (10) is achieved on. Moreover, the minimum in (10) occurs on some specific type and there is a polynomial number of types, so the set of metrics that receives the minimum of each type dominates the minimization and we can merge only a polynomial number of metrics to get the universal metric, which then follows to be optimal.

C. Finite state metrics

In this section we define a family of metrics that can be calculated using a finite state machine. This family was used in [8], where it was proved that the code length of the conditional Lempel-Ziv algorithm can be used as a universal decoding metric for finite state channels. For each n, let $\mathcal{S}_n$ be the state space with $|\mathcal{S}_n|$ states. A state machine is defined by the next state function $g:\mathcal{X}\times\mathcal{Y}\times\mathcal{S}_n:\to\mathcal{S}_n$, and $q:\mathcal{X}\times\mathcal{Y}\times\mathcal{S}_n:\to\mathbb{R}$ the output function. The first state is s_0. The metric $m(\mathbf{x}^n,\mathbf{y}^n)$ is computed as follows: Let $s_i=g(s_{i-1},x_i,y_i)$ for $i=1...n$, Then: $m(\mathbf{x}^n,\mathbf{y}^n)=\sum_{i=1}^n q(s_i,x_i,y_i)$.

The number of next state functions is: $|\mathcal{S}_n|^{|\mathcal{S}_n|\cdot|\mathcal{X}|\cdot|\mathcal{Y}|}$. Once the next state is given, only the type ($\mathbf{x},\mathbf{y}$ and the state) determines the metric value. In particular, the number of metrics that minimizes the pairwise error in (10) can be bounded by $B_n=|\mathcal{S}_n|^{|\mathcal{S}_n|\cdot|\mathcal{X}|\cdot|\mathcal{Y}|}\cdot(n+1)^{|\mathcal{X}|\cdot|\mathcal{Y}|\cdot|\mathcal{S}_n|}$. Let $|\mathcal{S}_n|=n^\alpha$. It is easy to see that $\lim_{n\to\infty}\frac{1}{n}\log\left(B_n\right)=0$ for $\alpha<1$ and $\lim_{n\to\infty}\frac{1}{n}\log\left(B_n\right)>0$ for $\alpha\geq 1$, so when $\alpha<1$ this family admits an optimal universal decoder.

When $|\mathcal{S}_n|=n$, *i.e.*$\alpha=1$, there exists a next state function such that the state in every different time is different, and in particular, given the sent and received words, we can construct a metric that gives maximal value to these two words and smaller metric value to other codewords. This means that the minimum in (10) is always $Q(\mathbf{x})$ and our decoder choose the codeword according to the a priori information, independently of $\mathbf{y}$. Clearly, this decoder is useless and not optimal.

Notice that metrics of different "next state function" admits no ties in general. Thus the type based universal decoder of [1] is not optimal, while our decoder, described above, is optimal for this family of finite state metrics with $\alpha<1$.

D. k^{th} order Markov metrics

A k^{th} order Markov metric is a special case of a finite state metric where the state contains the last k samples of $\mathbf{x},\mathbf{y}$. This class of metrics is interesting as it can be used to approximate several other interesting metrics, *e.g.*, metrics of stationary ergodic channels and finite state channels with hidden channel state. It is interesting to examine the possible Markov order (as a function of n) that allows universal decoding. Denote the order by k_n. The number of states is then $|\mathcal{S}_n|=\left(|\mathcal{X}|\cdot|\mathcal{Y}|\right)^{k_n}$. Universal decoding is possible as long as $|\mathcal{S}_n|=n^\alpha,\alpha<1$. Taking the $\log$ of both sides, we have: $k_n<\frac{\log(n)}{\log(|\mathcal{X}|\cdot|\mathcal{Y}|)}$.

VI. Summary and conclusions

A general universal decoder for a family of decoding metrics is described. This decoder generalizes several known universal decoders. The decoder, termed Generalized Minimum Error Test, uses the minimum error principle. We would like to explore the usability of the minimum error principle in other areas, including universal source coding with side information. Some further research regarding the infinite (abstract) alphabet is needed. Also, extension of the results here to network cases and to channels with feedback is also investigated in on-going research.

VII. Acknowledgment

Interesting discussions with Neri Merhav are acknowledged with thanks. Comments by the anonymous referees are greatly appreciated.

Appendix A
Rate functions

In order to analyze the performance of a metric decoder, we would need somehow to "normalize" the metric in order to remove its redundancy, *i.e.*, the performance depends only on the order induced by the metric and not the specific value of the metric. However, for rate function (which we will define here) the value of the metric provides a hint to the performance of decoding with this metric. Throughout we will refer to a distribution $Q_n(\mathbf{x})$ as the prior and we'll assume it to be fixed. Through the whole treatment above, the output sequence $\mathbf{y}$ always held fixed.

Definition 8.

1) A function $R : \mathcal{X}_n \to \mathbb{R} \cup \{-\infty\}$ is called **rate function** over the prior $Q_n(\mathbf{x})$ if:
$$Pr\{R(\mathbf{X}) \geq t\} = Q_n(R(\mathbf{X}) \geq t) \leq 2^{-nt}. \tag{20}$$
2) The function R is an **asymptotic rate function** if there exists a sequence λ_n such that $\lambda_n \underset{n\to\infty}{\to} 0$ and
$$Q_n(R(\mathbf{X}) \geq t) \leq 2^{-n(t-\lambda_n)}. \tag{21}$$
3) Given any function R (not necessarily a rate function) define:
$$\Omega_R(\mathbf{x}) = -\frac{1}{n}\log\left(Q_n(R(\mathbf{X}) \geq R(\mathbf{x}))\right). \tag{22}$$
$\Omega_R(\mathbf{x})$ is the **canonical rate function** associated with the order R.

Proposition 3.
(a) The function $\Omega_R(\boldsymbol{x})$ preserves the order induced by R on $\boldsymbol{x}$. That is, $R(\boldsymbol{x}_1) \leq R(\boldsymbol{x}_2)$ implies $\Omega_R(\boldsymbol{x}_1) \leq \Omega_R(\boldsymbol{x}_2)$. If $Q_n(\boldsymbol{x}_2) > 0$, then $R(\boldsymbol{x}_1) < R(\boldsymbol{x}_2)$ implies $\Omega_R(\boldsymbol{x}_1) < \Omega_R(\boldsymbol{x}_2)$.
(b) $\Omega_R(\boldsymbol{x})$ is a rate function.
(c) For any rate function R and any $\boldsymbol{x}$: $R(\boldsymbol{x}) \leq \Omega_R(\boldsymbol{x})$.

Proof:
(a) Let $f(t) = Q_n(R(\mathbf{X}) \geq t)$. $f(t)$ is a decreasing function. $\Omega_R(\mathbf{x}) = -\frac{1}{n}\log(f(R(\mathbf{x})))$ and (a) follows.
(b) Notice that since $\Omega_R(\mathbf{x})$ preserves the order in R, then $Q_n(\Omega_R(\mathbf{X}) \geq \Omega_R(\mathbf{x})) = Q_n(R(\mathbf{X}) \geq R(\mathbf{x})) = 2^{-n\cdot\Omega_R(\mathbf{x})}$ and the condition (20) is met with equality. For other t-s that are not equal to any $\Omega_R(\mathbf{x})$ the condition can be easily verified by noting that $Q_n(\Omega_R(\mathbf{X}) \geq t) = 2^{-n\cdot\Omega_R(\mathbf{x})}$ for some $\mathbf{x}$ such that $t < \Omega_R(\mathbf{x})$.
(c) If R is a rate function then
$$\begin{aligned} 2^{-n\cdot\Omega_R(\mathbf{x})} &= Q_n(\Omega_R(\mathbf{X}) \geq \Omega_R(\mathbf{x})) \\ &= Q_n(R(\mathbf{X}) \geq R(\mathbf{x})) \\ &\leq 2^{-n\cdot R(\mathbf{x})}. \end{aligned}$$
■

Example 3. *Let $R(\boldsymbol{x}) = \frac{1}{n}\log\left(\frac{P_n(\boldsymbol{x})}{Q_n(\boldsymbol{x})}\right)$ where $P_n(\boldsymbol{x})$ is a probability distribution, then R is a rate function. To see this, notice that by using Markov's inequality*
$$\begin{aligned} Q_n(R(\boldsymbol{X}) \geq t) &= Q_n(2^{n\cdot R(\boldsymbol{X})} \geq 2^{n\cdot t}) \\ &\leq \mathbb{E}_{Q_n(\boldsymbol{x})}\left(2^{n\cdot R(\boldsymbol{X})}\right)\cdot 2^{-n\cdot t} \\ &= \mathbb{E}_{Q_n(\boldsymbol{x})}\left(\frac{P_n(\boldsymbol{x})}{Q_n(\boldsymbol{x})}\right)\cdot 2^{-n\cdot t} \\ &= \sum_{\boldsymbol{x}} Q_n(\boldsymbol{x})\cdot\frac{P_n(\boldsymbol{x})}{Q_n(\boldsymbol{x})}\cdot 2^{-n\cdot t} = 2^{-n\cdot t}. \end{aligned}$$

Definition 9. We say that the function $R : \mathcal{X}_n \to \mathbb{R} \cup \{-\infty\}$ is an **asymptotically tight rate function** over the prior $Q_n(\mathbf{x})$ if there exists a sequence λ_n such that $\lambda_n \underset{n\to\infty}{\to} 0$ and

$$2^{-n(R(\mathbf{x})+\lambda_n)} \leq Q_n(R(\mathbf{X}) \geq R(\mathbf{x})) \leq 2^{-n(R(\mathbf{x})-\lambda_n)}. \tag{23}$$

Notice that for any function R, $\Omega_R(\mathbf{x})$ is a trivially asymptotically tight rate function (even not asymptotically). The following proposition gives sufficient conditions for any function R to be an asymptotically tight rate function.

Proposition 4. *For any function R we have:*

$$2^{-n\cdot\Omega_R(\boldsymbol{x})} \leq \mathbb{E}_{Q_n(\boldsymbol{x})}\left(2^{n\cdot R(\boldsymbol{x})}\right)\cdot 2^{-n\cdot R(\boldsymbol{x})}. \tag{24}$$

In particular, if

$$2^{-n\cdot(R(\boldsymbol{x})+\lambda_n)} \leq 2^{-n\cdot\Omega_R(\boldsymbol{x})} \tag{25}$$

and

$$\mathbb{E}_{Q_n(\boldsymbol{x})}\left(2^{n\cdot R(\boldsymbol{x})}\right) \leq 2^{n\cdot\chi_n} \tag{26}$$

with $\lambda_n, \chi_n \underset{n\to\infty}{\to} 0$, then R is an asymptotically tight rate function and $\mathbb{E}_{Q_n(\boldsymbol{x})}\left(2^{n\cdot\Omega_R(\boldsymbol{X})}\right)$ is also sub-exponential.

Proof: By Markov's inequality:

$$\begin{aligned}
2^{-n\cdot\Omega_R(\mathbf{x})} =& Q_n(R(\mathbf{X}) \geq R(\mathbf{x})) \\
=& Q_n(2^{n\cdot R(\mathbf{X})} \geq 2^{n\cdot R(\mathbf{x})}) \\
\leq& \mathbb{E}_{Q_n(\mathbf{x})}\left(2^{n\cdot R(\mathbf{X})}\right)\cdot 2^{-n\cdot R(\mathbf{x})}.
\end{aligned}$$

Combining this with (25), (26): $2^{-n\cdot(R(\mathbf{x})+\lambda_n)} \leq 2^{-n\cdot\Omega_R(\mathbf{x})} \leq 2^{-n\cdot(R(\mathbf{x})-\chi_n)}$. In other words, R is asymptotically tight. Observing that (25) implies that $2^{n\cdot(R(\mathbf{x})+\lambda_n)} \geq 2^{n\cdot\Omega_R(\mathbf{x})}$, and taking $\mathbb{E}_{Q_n(\mathbf{x})}(\cdot)$ of that we get $\mathbb{E}\left(2^{n\cdot\Omega_R(\mathbf{x})}\right) \leq 2^{n\cdot(\chi_n+\lambda_n)}$. ■

Recall from definition 7 that the prior $Q_n(\mathbf{x})$ is **normal** if $\mathbb{E}_{Q_n(\mathbf{x})}\left(2^{n\cdot\Omega_R(\mathbf{x})}\right)$ is subexponential for all $R(\mathbf{x})$. The following proposition bounds $\mathbb{E}_{Q_n(\mathbf{x})}\left(2^{n\cdot R(\mathbf{x})}\right)$ for any rate function $R(\mathbf{x})$ and provides a condition on the prior $Q_n(\mathbf{x})$ being normal.

Proposition 5.
(a) Suppose that there exists some B_n that bounds the rate function R from above, i.e., *$Pr\{nR(\boldsymbol{X}) \geq B_n\} = 0$. Then*

$$\mathbb{E}_{Q_n(\boldsymbol{x})}\left(2^{nR(\boldsymbol{X})}\right) \leq 1 + \ln(2)\cdot B_n. \tag{27}$$

(b) If $\mathcal{X}_n$ is discrete, $\chi_n = -\log\left(\min_{\mathbf{x}:Q_n(\boldsymbol{x})\neq 0} Q_n(\boldsymbol{x})\right)$ is a bound on any rate function. In particular, if χ_n is subexponential then the prior $Q_n(\boldsymbol{x})$ is normal in the sense of definition 7.

Proof: (a) By [15, Lemma 5.5.1], for a non-negative R.V. Z it holds that $\mathbb{E}(Z) = \int_0^\infty Pr\{Z \geq z\}\, dz$.

$$\begin{aligned}
\mathbb{E}\left(2^{nR(\mathbf{X})}\right) &= \int_0^\infty Pr\left\{2^{nR(\mathbf{X})} \geq t\right\} dt \\
&= \int_0^\infty Pr\{nR(\mathbf{X}) \geq \log(t)\}\, dt \\
&\overset{(1)}{\leq} 1 + \int_1^{2^{B_n}} Pr\{nR(\mathbf{X}) \geq \log(t)\}\, dt \\
&= 1 + \int_1^{2^{B_n}} Pr\left\{R(\mathbf{X}) \geq \frac{\log(t)}{n}\right\} dt \\
&\overset{(2)}{\leq} 1 + \int_1^{2^{B_n}} 2^{-n\cdot\frac{\log(t)}{n}} dt \\
&= 1 + \int_1^{2^{B_n}} \frac{dt}{t} \\
&= 1 + \ln(t)\big|_1^{2^{B_n}} = 1 + \ln(2)\cdot B_n.
\end{aligned}$$

Inequality (1) follows because $Pr\{\cdot\} \leq 1$ and $Pr\left\{nR(\mathbf{X}) \geq \log(2^{B_n})\right\} = 0$ and (2) because R is a rate function.

(b) For any $\mathbf{x}$ with $Q_n(\mathbf{x}) > 0$ we have:

$$\begin{aligned} n \cdot R(\mathbf{x}) &\leq n \cdot \Omega_R(\mathbf{x}) \\ &= -\log\left(Q_n(R(\mathbf{X}) \geq R(\mathbf{x}))\right) \\ &\leq -\log\left(Q_n(\mathbf{x})\right) \\ &\leq -\log\left(\min_{\mathbf{x}:Q_n(\mathbf{x})\neq 0} Q_n(\mathbf{x})\right). \end{aligned}$$

■

REFERENCES

[1] N. Merhav, "Universal decoding for arbitrary channels relative to a given class of decoding metrics," *IEEE Transactions on Information Theory*, vol. 59, no. 9, pp. 5566–5576, 2013.

[2] N. Elkayam and M. Feder. (2014) Exact evaluation of the random coding error probability and error exponent. [Online]. Available: http://www.eng.tau.ac.il/~elkayam/ErrorExponent.pdf

[3] V. D. Goppa, "Nonprobabilistic mutual information without memory," *Probl. Contr. Inform. Theory*, vol. 4, pp. 97–102, 1975.

[4] W. Huleihel and N. Merhav, "Universal decoding for gaussian intersymbol interference channels," *CoRR*, vol. abs/1403.3786, 2014.

[5] A. Lapidoth and E. Telatar, "Gaussian isi channels and the generalized likelihood ratio test," in *Information Theory, 2000. Proceedings. IEEE International Symposium on*. IEEE, 2000, p. 460.

[6] L. Farkas, "Blind decoding of linear gaussian channels with isi, capacity, error exponent, universality," *arXiv preprint arXiv:0801.0540*, 2008.

[7] N. Merhav, "Universal decoding for memoryless gaussian channels with a deterministic interference," *Information Theory, IEEE Transactions on*, vol. 39, no. 4, pp. 1261–1269, 1993.

[8] J. Ziv, "Universal decoding for finite-state channels," *Information Theory, IEEE Transactions on*, vol. 31, no. 4, pp. 453–460, 1985.

[9] A. Lapidoth and J. Ziv, "On the universality of the lz-based decoding algorithm," *Information Theory, IEEE Transactions on*, vol. 44, no. 5, pp. 1746–1755, 1998.

[10] Y. Lomnitz and M. Feder, "Universal communication over arbitrarily varying channels," *IEEE Transactions on Information Theory*, vol. 59, no. 6, pp. 3720–3752, 2013.

[11] ——, "Universal communication - part i: Modulo additive channels," *IEEE Transactions on Information Theory*, vol. 59, no. 9, pp. 5488–5510, 2013.

[12] M. Feder and A. Lapidoth, "Universal decoding for channels with memory," *Information Theory, IEEE Transactions on*, vol. 44, no. 5, pp. 1726–1745, 1998.

[13] Y. Lomnitz and M. Feder, "Communication over individual channels–a general framework," *arXiv preprint arXiv:1203.1406*, 2012.

[14] N. Shulman, "Communication over an unknown channel via common broadcasting," Ph.D. dissertation, Tel Aviv University, 2003.

[15] H. Koga *et al.*, *Information-spectrum methods in information theory*. Springer, 2002, vol. 50.